

UK Patent Application GB 2 177 877 A

(43) Application published 28 Jan 1987

(21) Application No 8517735 ·

(51) INT CL⁴
H04L 11/20 ·

(22) Date of filing 13 Jul 1985 ·

(52) Domestic classification (Edition I)
H4P PS ·

(71) Applicants
STC plc.

(56) Documents cited
EP A1 0130431
EP A1 0016477 ·
WO A1 84/00268 ·

(Incorporated in United Kingdom), ·

(58) Field of search
H4P ·
H4K ·
Selected US specifications from IPC sub-class H04L ·

190 Strand, London WC2R 1DU

British Telecommunications plc.

(Incorporated in United Kingdom),

81 Newgate Street, London EC1A 7AJ

(74) Agent and/or Address for Service

(72) Inventors
Geoffrey Harland,
Ronald Frank Rous,
Alistair Russell McKenzie,
Peter James Gee .

S. R. Capsey,
STC Patents, Edinburgh Way, Harlow, Essex CM20 2SH ·

(54) Packet systems

(57) In a packet switching system, e.g. a local area network, in which synchronous data such as voice or video traffic is handled, it is essential that measures be taken to deal with variable delays in transmission which could lead to loss of bytes which suffer excessive delay. In the present system, the reconstruction of a bit stream from received packets is controlled by time tagging of packets which taggings are noted at the receiving end, or by the use of a fixed or adaptive delay.

Excessive packet delays, which cause loss of bytes, results in the lost bytes being substituted by the repetition of the previous byte, a zero byte, or one representing half the value of the previous byte.

GB 2 177 877 A

2177877

112

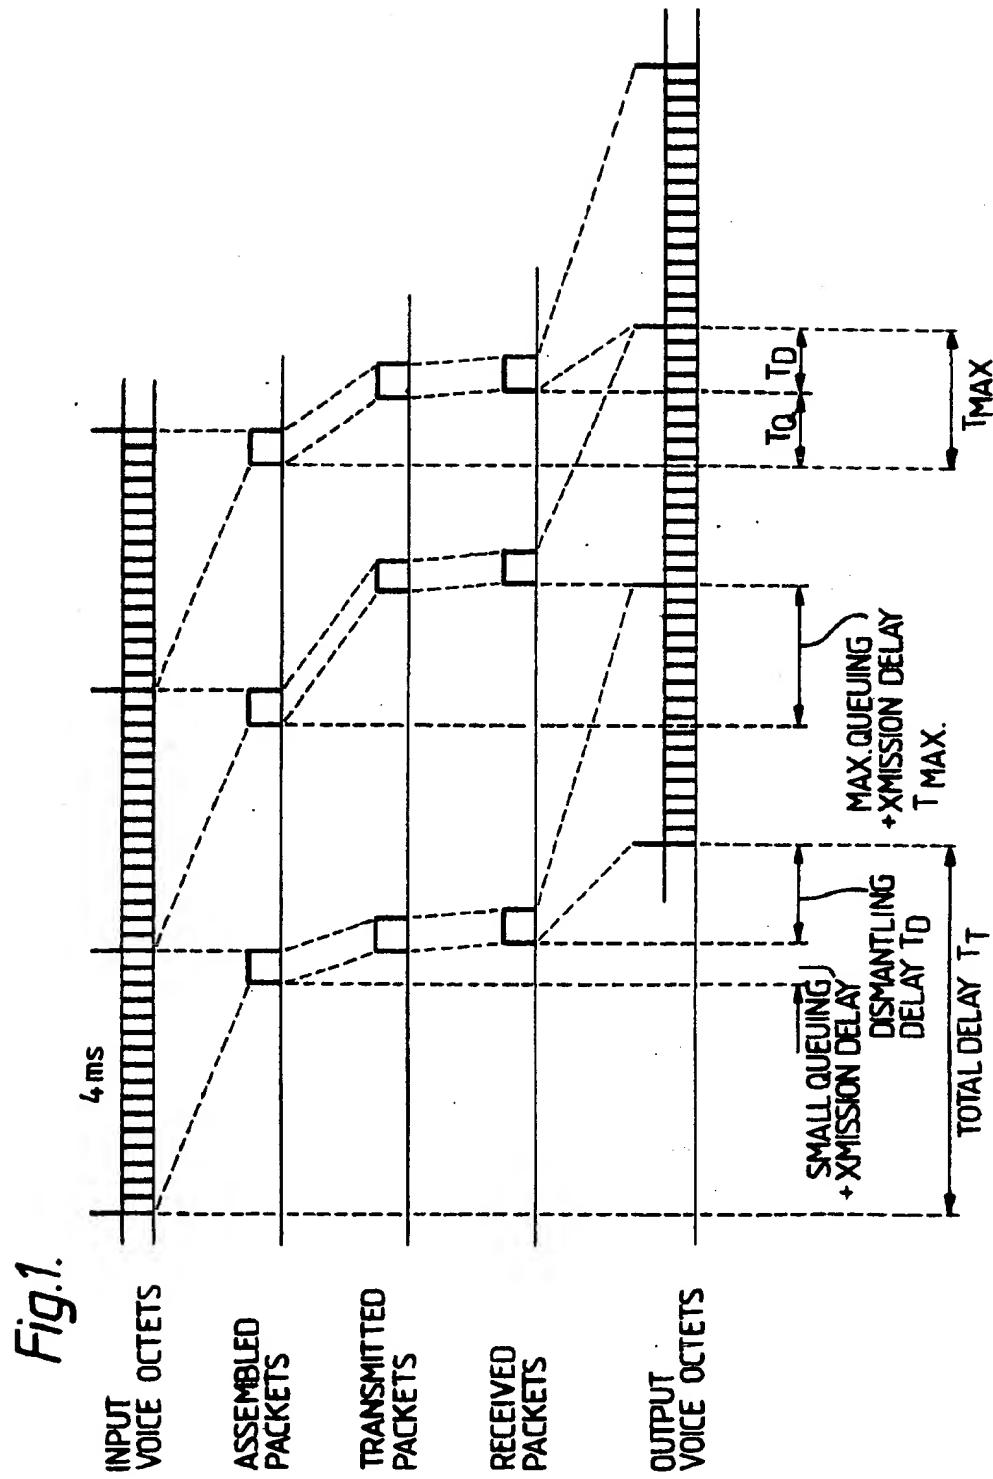
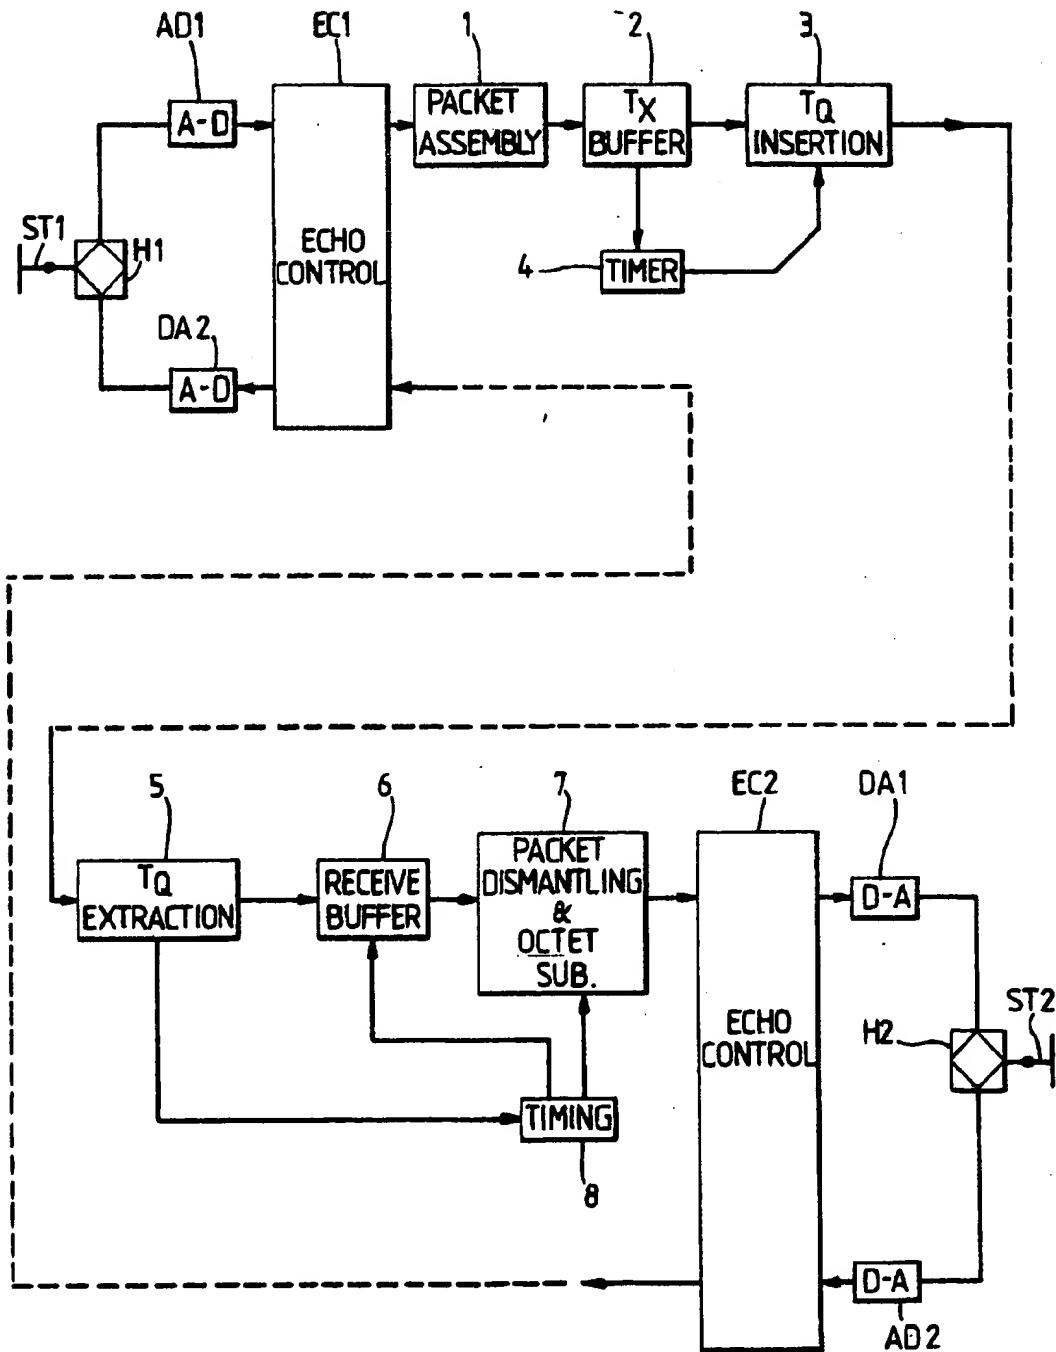


Fig.1.

2177877

2/2

Fig.2.



SPECIFICATION
Packet systems

This invention relates to packet switching systems, especially, but not solely, to such a system when used in a Local Area Network (LAN).

In handling voice and other synchronous services in a packet switch or packet switched network, it is necessary to ensure that the probability that delayed packets arrive too late to reconstruct the synchronous bit stream is very small. To ensure this, it is necessary for the dismantlement of minimally delayed packets to be delayed such that the synchronous bit stream timing is matched to the maximally delayed packets.

According to the invention, there is provided a packet switching system, in which synchronous data such as voice or video data is handled, each packet as transmitted including a number of data bytes accompanied by additional control or signalling bytes, in which the parameters of the system are such that packets are liable to be subjected to different delays in transmission between two points, in which reconstruction of a synchronous bit stream from the received packets involves the dismantlement of each received packet and the assembly of the data in the packets into a bit stream forming the received data, in which the timing of said dismantlement and assembly is controlled by:

- (a) time tagging of the packets, whereby the first packet which is transmitted includes an indication of the time at which that packet is sent; or
- (b) the introduction at the receiving end of a fixed or an adaptive delay to some or all of the received packets, and in which if excessive packet delays result in bytes of a packet being lost, additional bytes are inserted to complete the assembled bit stream, each such additional byte being selected in a preset manner, or in one of a plurality of preset manners.

An embodiment of the invention will now be described with reference to the accompanying drawings, in which Fig. 1 is a timing diagram relating to voice packet delays in a packet switched system, while Fig. 2 is a simplified schematic of a packet system embodying the invention.

The following description relates to the application of the invention to a LAN, but it will be understood that the idea is easily extrapolatable to cover packet switched networks with tandemed transmission links.

In the arrangement of Fig. 2, we see a terminal ST1 at one end coupled to the system via a hybrid (or its electronic equivalent) H1, with another terminal ST2 and hybrid (or equivalent) H2 at the other end. Voice data to be handled passes from ST1 via H1 to an analogue to digital converter AD1 from which it is applied to echo control circuitry EC1, if such circuitry is needed.

From the block EC1, if present, the digital data is applied to a packet assembly circuit 1, from which the packets pass via a transmission buffer 2 to a T_o insertion block 3, whose function is indicated below.

Associated with these last two blocks is a timer 4.

The output of the block 3 is applied to the shared transmission medium on the GO side, which medium, e.g. a coaxial cable or an optical fibre cable, is shown in dashed lines. The transmission medium in most cases include contention resolution and bit encoding/decoding arrangements (not shown).

At the other end, the packets reach the T_o extraction block 5, from which the packets pass via a receive buffer 6 to a packet dismantling and octet insertion block 7. Blocks 6 and 7 function under control of a timing unit 8, controlled from the T_o extraction block 5.

The output from the block 7 goes to further echo control circuitry EC2 if such circuitry is needed. The output of EC2, if present, goes via a digital to analogue converter DA1 to the hybrid H2, from which it reaches the terminal ST2.

In the reverse direction, the transmission path is via an analogue to digital converter AD2 and the block EC2, from which it passes to EC1, in the same manner as the packets from ST1 pass from EC1 to EC2. At the first end, the data dismantled from the received packets pass via a digital to analogue converter DA2 to the hybrid H1 and therefrom to the terminal ST1.

In the following description, we refer both to Fig. 1, which relates to voice packet delays in a packet switch (which is generally similar to what applies to a LAN), and Fig. 2 which is the system block schematic referred to above.

Input voice octets, i.e. eight-bit bytes, are assembled into packets in the packet assembler 1 during 4ms (for say 32-octet packets) and queued in the transmit buffer 2 for access to the transmission medium. Transmission occurs after a variable delay, dependent on network activity, as access delay increases with traffic. Packet reception suffers only a small additional fixed delay due to propagation time.

Packets which suffer the maximum allowable queueing delay should be dismantled in the packet dismantling and octet substitution block 7, immediately on receipt. Dismantling of packets which experience smaller queueing delays must, however, be further delayed so that the resultant bit stream, which is derived from the less delayed and the maximally delayed packets, is contiguous, the bit stream due to the maximally delayed packets, and that due to the less delayed packets, is contiguous, as shown in Fig. 1. The problem is for the receiving node to determine the dismantling delay (T_o) required, from a set of randomly jittering packets.

The simplest solution to the problem of differing delays is to delay dismantling the first voice packet of a call by a period T_{max} . This does the job but introduces an additional delay to all packets equal to the queueing delay T_o of the first packet. This delay T_o can take any value from zero (ignoring propagation delay) to T_{max} or with small probability even greater. Its average value varies with the LAN loading, e.g. at 0.7 loading, average T_o is approximately $T_{max}/6$, for a condition in which 1 in 10^3 packets are delayed by periods each greater

than T_{max} . This assumes that packet length is constant. The additional delay reduces the probability of octet loss, on average, but cannot be relied upon to help the individual call. Additional space in the buffer 6 at the receive node is necessary to provide this. The network has to be so dimensioned that T_{max} is only half the maximum allowable delay.

5 A slightly more complex procedure uses an adaptive arrangement. Little is to be gained by delaying dismantling the first voice packet by only $T_{max} - T_0$ average and then adapting upwards later, if necessary, since the difference between this and T_{max} is probably not worth the trouble. However, it 10 might be worthwhile to start (say) $T_{max} - T_0$ average and then adapt downwards by preventing the dismantling delay from ever exceeding T_{max} . This would involve the insertion of dummy octets at the block 7, where necessary, during the call. The 15 arrangement quickly adapts to the average queueing delay during the first second (say) of the call, with little disturbance after that.

Another, and perhaps, better approach is to use the fact that, in most LAN systems, the transmitting 20 node has (in principle) knowledge of the instantaneous queueing delay. Thus the first voice packet of a call could be used to transmit the value of T_0 (for that packet) to the receive node, this value being inserted at block 3 under control of the timer 25 4. At the receiving node, the value T_0 is subtracted from T_{max} in the block 5 to give the required initial value of T_0 which is applied to the buffer 6. This ensures that, for the remainder of the call, $T_0 + T_0$ always equal T_{max} . This assumes that the two 30 nodes are synchronous at the 64 kbit/s rate.

In larger packet switched networks, each switching node transmits the value of the cumulative queueing delay, i.e. each node adds its queueing delay to the value sent from the preceding 35 node. If the packet dismantling node does not receive a delay value, due to not all parts of the complete packet connection having this capability or due to a system fault, it could use the fixed or adaptive dismantling delay technique as a fall back 40 arrangement.

Treatment of Excessively Delayed Packets

Packets delayed more than the maximum allowed for by the dismantling process discussed above 45 cause voice or other synchronous service octets not to be available for insertion in the reconstructed synchronous bit stream. The number of unavailable octets depends on the extent to which the packet is delayed relative to the reconstructed bit stream. 50 Thus a delay excess of up to 125 μs causes loss of only one octet at 64 kbit/s. The probability of losing two or more octets due to excess delay is generally much lower than the probability of losing only one.

When octets are lost due to excessive delay, it is 55 necessary to put something in their place in the reconstructed bit stream. The substitution effected in the block 7, is so chosen as to minimise the impact on the service supported, e.g. to minimise clicks in voice. For some services, the preferred 60 substitution is the repetition of the previous octet; in

others it is the insertion of a 'zero' octet. In some cases, it may be preferred to substitute an octet having half the value of the previous octet, so that the value declines to zero if several octets are 65 missing.

Synchronous Operation

Plesiochronous slip can be a severe problem for synchronous services and particularly for 2 Mbit/s video. Unless quite complex recovery mechanisms are worked out between LAN and video codec designers (and implemented), frequent picture break-up occurs. For example, with 8000 2 Mbit/s frames per second and 1 in 10^4 clock frequency 70 difference, a 2 Mbit/s frame could be lost every 1.25 seconds. Without suitable recovery mechanisms, the picture quality becomes continuously unacceptable, since normal picture update (even with no movement) is very slow. Even use of the fast update recovery mechanism under LAN control, which takes 60 ms, might be disturbing if it occurred as frequently as once every few seconds.

In view of this, it is worth going to some trouble to provide true synchronism, where desirable. One 75 solution is to use some form of synchronous LAN but other approaches are possible.

The clocks used for (say) 20 Mbit/s transmission on the LAN do not need to be synchronised, for the transmission of data (including voice and video) to 80 be truly synchronous. The transmitted data include their own timing and it is generally possible to recover this at the receiving node. For example, 32 voice octets received at 4 ms intervals, on average, should enable a 64 kHz clock to be synchronised.

100 Jitter due to 20 Mbit/s "granularity" should be insignificant at 64 kbit/s and lower data rates, and even at 2 Mbit/s may not be serious. Jitter at 250 Hz, due to variable queueing times, should also be acceptable. If necessary, however, the value of T_0 , 105 see above, can be transmitted with every data (e.g. voice) packet, to enable time (relative to the transmitting node) to be determined exactly in block 7 at the receiving node, under control of timing block 8.

110 The master clock should be the 2 MHz clock from ISDN (Integrated Services Digital Network), on calls to/from that network. Similarly, digital tie line gateways should provide the master clock, on relevant calls. Internal synchronous calls could 115 probably rely on mutual synchronisation, by suitable control loop design.

The above arrangements make it possible to provide truly synchronous services across a LAN. This eases the problems of meeting voice 120 performance constraints, allows much simpler recovery mechanisms for 2 Mbit/s video, and enables synchronous data services to be provided. It also reduces buffering requirements and simplifies buffer control.

125

CLAIMS

1. A packet switching system, in which synchronous data such as voice or video data is handled, each packet as transmitted including a 130 number of data bytes accompanied by additional

control or signalling bytes, in which the parameters of the system are such that packets are liable to be subjected to different delays in transmission between two points, in which reconstruction of a synchronous bit stream from the received packets involves the dismantlement of each received packet and the assembly of the data in the packets into a bit stream forming the received data, in which the timing of said dismantlement and assembly is controlled by:

- (a) time tagging of the packets, whereby the first packet which is transmitted includes an indication of the time at which that packet is sent; or
- (b) the introduction at the receiving end of a fixed or an adaptive delay to some at least of the received packets,

and in which if excessive packet delays result in bytes of a packet being lost, additional bytes are inserted to complete the assembled bit stream, each such additional byte being selected in a preset manner, or in one of a plurality of preset manners.

2. A system as claimed in claim 1, in which the additional byte which is included in a bit stream when a byte has been lost is either (a) a repeat of the previous byte, (b) a byte representing half the value of the previous byte, or (c) a zero amplitude byte.

3. A system as claimed in claim 1, in which the packet switching system does not inherently provide synchronism between transmitting and receiving modes and in which all packets for synchronous services are time tagged, to facilitate synchronisation of the reconstructed bit stream with the input bit stream.

4. A packet switching system substantially as described with reference to the accompanying drawings.

Printed for Her Majesty's Stationery Office by Courier Press, Leamington Spa, 1/1987. Demand No. 8817356.
Published by the Patent Office, 25 Southampton Buildings, London, WC2A 1AY, from which copies may be obtained.